

# Proportional Fair Rate Allocation for Private Shared Networks

Saman Feghhi, Douglas J. Leith, Mohammad Karzand  
 School of Computer Science and Statistics, Trinity College Dublin  
 Email: {feghhis,doug.leith,karzandm}@tcd.ie

**Abstract**—In this paper, we consider fair privacy in a shared network subject to traffic analysis attacks by an eavesdropper. We initiate the study of the joint trade-off between privacy, throughput and delay in such a shared network as a utility fairness problem and derive the proportional fair rate allocation for networks of flows subject to privacy constraints and delay deadlines.

## I. INTRODUCTION

Privacy in packet switched networks has attracted increasing interest in the research community over the last decade. Traditionally encryption techniques have been used to provide privacy by hiding the content of communication messages from an eavesdropper. However, more sophisticated attacks have been developed in recent years that use characteristics other than the content of the packets to gain information about the content of transmissions. For example, it has been demonstrated that packet length, count and the inter arrival times between packets can be used to infer with high accuracy the identity of web pages being browsed [6], video being watched and the content of encrypted voice calls [16], [17]. For attacks which make use of packet length information the obvious defence is to pad/fragment packets to be of fixed size. However, this simple defence is insufficient to protect against timing only attacks, since these make no use of packet size information. Note that timing only attacks can be powerful e.g. in [6], [7] it is demonstrated that the web site being browsed can be successfully inferred with greater than 90% accuracy using timing information alone.

Timing attacks make use of the properties of the packet stream, rather than of individual packets, and defence against such attacks therefore requires use of traffic shaping to modify the timing pattern of the stream of packets transmitted over the network such that it becomes hard for an eavesdropper to learn about the original message. Such traffic shaping can be achieved by inserting dummy packets into the packet stream to mask idle periods, by delaying/buffering packets to modify their timing and of course by dropping packets. In practice this shaping might be performed by, for example, end hosts or by a VPN gateway.

However, such privacy enhancing traffic shaping can impose a cost on the user and on the network. For example, insertion of dummy packets increases the load on the network while buffering user packets for longer increases delay but can reduce the need for dummy packets to ensure privacy. Importantly, for users sharing common a network path this




Fig. 1: Schematic illustrating simple two flow example.

creates a joint trade-off between their privacy, throughput and delay. This is illustrated in Figure 1 which shows two users whose traffic is shaped to enhance privacy and is then sent over a shared link which may be subject to eavesdropping. An increase in the rate of dummy packet transmissions by user 1 reduces the available bandwidth for user 2 on the shared link, and so to maintain privacy may require user 2 to reduce the rate at which useful (non-dummy) packets are sent, to buffer packets for longer (to disrupt timing without adding extra dummy packets) or to drop more packets (again to disrupt timing patterns). Alternatively, user 2 may choose to sacrifice privacy in order to avoid reducing throughput, increasing delay *etc.* That is, increased privacy for one user may come at the cost of decreased throughput, increased delay and/or reduce privacy for another user.

In this paper, we initiate the study of the joint trade-off between privacy, throughput and delay in a shared network as a utility fairness problem. We derive the proportional fair rate allocation for networks of flows subject to privacy constraints and delay deadlines. To the best of our knowledge this is the first study of *fair* privacy in a shared network subject to timing attacks.

## II. RELATED WORK

There exists a large body of research focused on fair rate allocation in shared networks but without consideration of privacy, see for example [5], [11], [1] and references therein. With regard to privacy, fairness has previously been considered in the context of mix networks [2] where the aim is to achieve unlinkability/anonymity. Mishra et al. [12] propose a proportional fair scheduling model that preserves source anonymity and in [13] investigate the trade-off between anonymity and quality of service. The focus in both papers is on hiding the identity of the users in a shared network. With regard to traffic analysis and associated defences, previous studies

have established that the packet size, count, transmission rate, inter-arrival times *etc.* can reveal significant information about a traffic flow, see e.g. [16], [4], [6], [7] and references therein. A number of defences have been proposed to counter these traffic analysis attacks, mostly addressing attacks that make use of packet size, count and the length of transmission *e.g.* they modify the traffic by padding the packets to be the same length or by adding dummy traffic to the end of the packet sequence to mask the packet count [8]. These defences are, however, ineffective against attacks that solely use timing information of network flows, see for example [6] and [7]. To the best of our knowledge none of this previous work on defences against traffic analysis attacks has considered fairness or the joint aspect of the trade-off between privacy, throughput and delay in a shared network.

### III. ATTACK MODEL & PRIVACY FOR A NETWORK FLOW

We consider a network path where time  $t \in \{0, 1, \dots\}$  is slotted and in slot  $t$  a new packet may arrive for transmission across the network. We let random variable  $X(t) = 1$  if a packet arrives in slot  $t$  and  $X(t) = 0$  otherwise. It is assumed that packets themselves contain no interesting information for an eavesdropper, which means that the packets are strongly encrypted and are of constant size. However, the sequence  $\mathbf{X} := \{X(t)\}$  of packet arrivals contains information that can be used to reveal the characteristics of the traffic flow. To protect this information, the arriving packet stream is passed through a traffic shaper before being transmitted across the network. We let random variable  $Y(t) = 1$  if a packet is transmitted across the path in slot  $t$  and  $Y(t) = 0$  otherwise.

Our attack model is that the sequence  $\mathbf{Y} := \{Y(t)\}$  of transmissions is observable by an adversary, but not the sequence  $\mathbf{X}$  of arrivals. Note that in general we expect that an eavesdropper listening to network transmissions does not observe output sequence  $\mathbf{Y}$  directly, but rather only after transformation by the MAC layer *i.e.* after buffering, scheduling delay *etc.* This will generally increase privacy (this follows from the data processing inequality, and more specifically *e.g.* passing through a queue increases entropy). Nevertheless, we take a conservative approach and assume a strong attacker who can perfectly invert these changes and so recover  $\mathbf{Y}$ .

We use the mutual information  $I(\mathbf{Y}; \mathbf{X})$  between sequences  $\mathbf{X}$  and  $\mathbf{Y}$  as our measure of privacy, similarly to [18] and others. Mutual information measures the capacity of the information channel between sequences  $\mathbf{X}$  and  $\mathbf{Y}$ . When  $I(\mathbf{Y}; \mathbf{X}) = 0$  the sequences are statistically independent and more generally mutual information captures the exposure to watermarking attacks, which can be thought of as a worst case situation in our setup, namely where the user sequence  $\mathbf{X}$  has a signature intentionally embedded within its timing pattern by an adversary (e.g. by a web site which is being downloaded). Recall that  $I(\mathbf{Y}; \mathbf{X}) = H(\mathbf{Y}) - H(\mathbf{Y}|\mathbf{X})$ , where  $H(\mathbf{Y})$  is the entropy of the output sequence and  $H(\mathbf{Y}|\mathbf{X})$  the conditional entropy between the output and input sequences. Privacy is maximised when  $I(\mathbf{Y}; \mathbf{X})$  is minimised (as already noted, when  $I(\mathbf{Y}; \mathbf{X}) = 0$ , the output sequence  $\mathbf{Y}$  is statistically independent from the input sequence  $\mathbf{X}$  and we say transmissions are fully private).

### IV. ON-OFF TRAFFIC SHAPING POLICY

Transmitting a packet in every slot regardless of the pattern of the arriving user traffic ensures that the transmitted packet sequence contains no information about the user traffic<sup>1</sup> and  $I(\mathbf{Y}; \mathbf{X}) = 0$ . Similarly, when no packets are transmitted,  $Y(t) = 0$ ,  $t = 1, 2, \dots$ , then also trivially  $I(\mathbf{Y}; \mathbf{X}) = 0$ . However, transmitting a packet in every slot means that when there is no information packet to send a dummy packet must be transmitted and so is clearly wasteful. And transmitting no packets at all is clearly private but not useful for communication. This motivates use of an on-off approach to transmission. That is, we specify an interval  $\tau$ . In the first slot of this interval we transmit a packet (sending a dummy packet if no information packets are available<sup>2</sup>), and over the remaining  $\tau - 1$  slots no packets are transmitted (any arriving information packets are buffered in a queue). That is,  $Y(t) = 1$  for  $t = 1$  and  $Y(t) = 0$  for  $t \in \{2, \dots, \tau\}$ . It is known from queueing theory that deterministic service minimises the queueing delay for a specified mean service rate [10], *i.e.* periodic service minimises delay. Conveniently, periodic service also means that the pattern of packet transmissions contains no information, *i.e.* the mutual information between the transmitted sequence  $\mathbf{Y}$  and the arrival process  $\mathbf{X}$  is zero,  $I(\mathbf{Y}; \mathbf{X}) = 0$  and so transmissions are fully private. Hence this traffic shaping approach is a minimum delay maximum privacy one. By adjusting the mean transmit rate via parameters  $g$  and  $\tau$  we can tune the trade-off between the number of dummy packets sent (which waste network bandwidth) and the buffering delay experienced by information packets.

This on-off traffic shaper can be modelled as a so-called Fixed Cycle Traffic Light (FCTL), first studied in [14] in the context of vehicular traffic. Consider cars arriving at a junction controlled by a FCTL. The light has two states which divides the total cycle into fixed length green ( $g$ ) and red ( $r$ ) cycles. In the red cycle, a new arrival enters a queue, waiting to cross the junction. When the light turns green, cars in the queue cross the junction one at a time until the cycle lasts or queue becomes empty. Cars arriving at the junction during a green cycle and finding the queue empty proceed to cross the junction. It is this latter characteristic which makes the FCTL different from a conventional queue with periodic service. In our setup the green cycle is of duration  $g = 1$  slot and the red cycle is of duration  $r = \tau - 1$  slots. The characteristics of a FCTL have been much studied, see for example the overview in [3], with estimates of the average queue length given in [14] and [15].

#### A. Delay

Suppose the input sequence  $\{X(t)\}$  is i.i.d with  $\mathbb{P}(X(t) = 1) = p$  and  $\mathbb{P}(X(t) = 0) = 1 - p$ . That is, the traffic

<sup>1</sup>Formally, suppose output sequence  $Y(t) = 1$ ,  $t = 1, 2, \dots$  *i.e.* a packet is transmitted in every slot. The entropy of the transmission at a single slot is  $H(Y(t)) = p \log p + (1 - p) \log(1 - p)$  where  $p = \text{Prob}(Y(t) = 1)$ . Since  $p = 1$  when a packet is always transmitted,  $H(Y(t)) = 0$ . The entropy of the sequence  $\{Y(t)\}$  satisfies  $H(\mathbf{Y}) \leq \sum_t H(Y(t)) = 0$  and since  $H(\mathbf{Y}) \geq 0$  it follows that the entropy of the sequence  $\{Y(t)\}$  is zero. Similarly,  $H(\mathbf{Y}|\mathbf{X}) = 0$  and so the mutual information  $I(\mathbf{Y}; \mathbf{X}) = 0$ .

<sup>2</sup>Note that it is also possible to adopt a partially private approach where, when no information packet is available to send, a dummy packet is transmitted with some probability less than one. In this case the mutual information  $I(\mathbf{Y}; \mathbf{X})$  will be non-zero, but can be controlled by adjusting the probability with which dummy packets are sent. However, we leave this more general case for future work.




Fig. 2: Comparison of Newell's (equations (3.7) and (5.6) in [14]) and Miller's estimates for the waiting time at a FCTL with simulation data. The simulation values are averaged over 1000 cycles with  $p = 0.3$  and  $\tau = 100$ . The data for  $g/\tau < p$  is excluded as the queue is unstable in this region.

arrivals form a Bernoulli process. Given that the time is slotted, following [14], the average waiting time  $w$  for each packet measured in time slots is

$$w = (\tau - g)(1 - p)^{-1}\tau^{-1}(E(q_x)/p + (\tau - g + 1)/2). \quad (1)$$

where  $E(q_x)$  denotes the length of queue right at the end of green (serving) period of the  $x$ th cycle and  $\tau$  is the length of each cycle *i.e.*  $\tau = r + g$ .

In order to calculate  $w$  we need to find the equilibrium distribution for  $q_x$  and evaluate  $E(q_x)$ . While the calculation of  $q_1$  and  $q_2$  is straightforward, the evaluation of  $q_3$  onwards quickly becomes tedious and the expressions cumbersome. Several attempts have therefore been made to derive accurate estimates of  $E(q_x)$  which are simpler in nature. Two estimates, that of Newell [14] and of Miller [15], are illustrated in Figure 2 and compared against numerical simulation data. In the sequel we use Miller's estimate due to its accuracy and simplicity. This is given by

$$E(q_x) \approx \max\left\{\frac{2p\tau - g}{2(g - p\tau)} \cdot (1 - p), 0\right\} \quad (2)$$

### B. Rate of Dummy Packet Transmissions

In addition to the delay introduced by traffic shaping we are also interested in the fraction  $d$  of slots expended on dummy packet transmissions (when no information packet is available to send at a slot in an on cycle). The information packet arrivals  $\nu$  during an off period are distributed as:

$$\text{Prob}(\nu = k) = \binom{T}{k} p^k (1 - p)^{T-k}, \quad (3)$$

and the associated probability generating function is  $K(z) = [(1 - p) + pz]^T$ . Regarding  $\Pi(z)$ , the probability generating function for the limiting distribution  $\pi_j$  the length of the queue at the start of the first slot of an on period, we have:

$$\Pi(z) = \frac{\sum_{j=0}^{g-1} \pi_j (z^g - z^j)}{z^g/K(z) - 1} = \frac{\sum_{j=0}^g \pi_j (z^g - z^j)}{z^g[(1 - p) + pz]^{-T} - 1}. \quad (4)$$

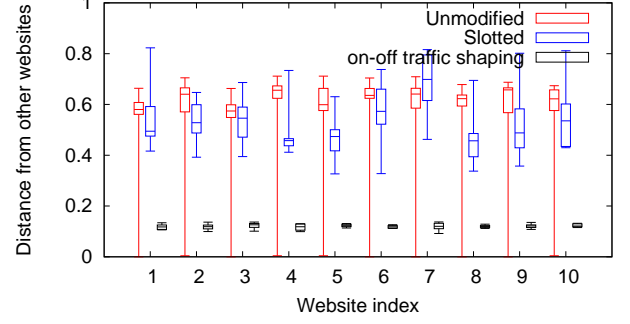


Fig. 3: Comparison of  $F$ -distance between different web traces for unmodified, slotted and queued+slotted scenarios. The slot size is 0.01s,  $g = 5$ ,  $\tau = 10$  and the window size for evaluating  $F$ -distance is 0.2.

We know  $\Pi(1) = 1$ , which means after using L'hopital's rule

$$d = \frac{1}{\tau} \sum_{j=0}^{g-1} \pi_j (g - j) = \frac{g}{\tau} - p, \quad (5)$$

the average number of dummy packets transmitted over cycle.

We will assume that  $g - p\tau > 0$ . This means that on average our service can accommodate all of the arrivals to the queue and the queue is stable. Stability can be seen by looking at the polynomial  $z^{T_{on}}[(1 - p) + pz]^{-T} - 1$  which, for stability, should have no zeros in or on the unit circle. Now,

$$z^g = [(1 - p) + pz]^\tau, \quad gz^g = p\tau[(1 - p) + pz]^{\tau-1} \quad (6)$$

Dividing the two equations, for  $|z| < 1$  we should have  $p\tau > g$ , but this is excluded when  $g - p\tau > 0$ .

### C. Example

In order to investigate the effectiveness of on-off traffic shaping on privacy, we conducted the following test: we collected packet timestamp traces from 10 different web sites. The traces are of approximately same length. We compared all of the samples with each other using Dynamic Time Warping and calculated their  $F$ -distance [6]. The  $F$ -distance, which is a value in  $[0, 1]$ , is a measure of the similarity between two timestamp traces and is used in [6] as the basis for a successful timing-only traffic analysis attack based in  $k$ -NN clustering. The  $F$ -distance is smaller when traces are similar and increases as they become more different. It can be seen in Figure 3 that the average distance between the traces for different web sites is relatively large, which would enable an attacker to distinguish between them.

We then manually divided time into slots (of 10ms duration) and adjusted the packet timings so that they are only sent at the beginning of a slot<sup>3</sup>. It can be seen in Figure 3 that time slotting decreases the average distance between packet traces, the decrease is relatively small and would still allow an attacker to distinguish between different the traces for different web sites.

<sup>3</sup>Note that using an actual time slotted tunnel adds an additional distortion due to network protocols. So by manually time slotting the time we are considering a slightly easier scenario for the attacker.




Fig. 4: Comparison between transmission duration and packet count for unmodified and on-off shaped web traces.  $g/\tau = 0.5$ .




Fig. 5: Comparison between transmission duration and packet count for unmodified and on-off shaped web traces.  $g/\tau = 0.01$ .

Finally we applied on-off traffic shaping and calculated the  $F$ -distance. It can be seen that there is now a considerable drop in the average distance, and also in the variance. This indicates that the modified traces are now much more similar, making it difficult for the attacker to distinguish between web sites.

As already noted, traffic shaping introduces a queueing delay and the transmission of additional dummy packets. Figure 4 shows the measured delay and number of dummy packets when  $g/\tau = 0.5$ . It can be seen that the delay is negligible since  $\tau$  is small but the number of dummy packets is almost 2 times that of the real traffic. As  $g/\tau$  is decreased, the number of dummy packets transmitted falls, but the delay increases. This is illustrated in Figure 5 for  $g/\tau = 0.01$ .

## V. PROPORTIONAL FAIR PRIVACY

We consider a shared network where each user applies traffic shaping before transmitting their traffic over the network, e.g. see Figure 1. Note that traffic shaping is applied individually to each user's traffic, which is required to prevent active traffic injection attacks (where an adversary sharing a queue with a user injects traffic in a way that reveals the presence/absence of a user packet in the queue [9]). In practice this setup might correspond to a VPN where the client shapes and encrypts arriving traffic before transmitting it across the Internet to a VPN gateway. As already noted, buffering and shaping may introduce delay (via buffering) and consume bandwidth (via dummy packet transmissions) and there is a joint trade-off between privacy, delay and throughput for the users sharing the network.

### A. Network Model

We consider a multi-user network serving a set of  $\mathcal{N}$  users and a set  $\mathcal{F}$  of flows. Each flow  $f \in \mathcal{F}$  has a source  $s_f \in \mathcal{N}$ , a mean offered load rate  $\mu_f$  and a mean delay deadline  $\sigma_f$ . The network uses a scheduler that can service the offered load provided the aggregate flow usage satisfies

$$\sum_{f \in \mathcal{F}} \frac{\mu_f}{\psi_f} \leq 1 \quad (7)$$

where  $\psi_f$  is the physical transmit rate for flow  $f$  and so  $\mu_f/\psi_f$  is the mean fraction of airtime used by flow  $f$  and for arrival rate  $p_f$  and dummy rate  $d_f$  we have  $\mu_f = p_f + d_f$ .

### B. Convexity of Traffic Shaping Delay

As already noted, on-off traffic shaping introduces additional delay. When on-time  $g = 1$  we have

$$E(q_x) \approx \max\left\{\frac{(2p - c)(1 - p)}{2(c - p)}, 0\right\} \quad (8)$$

and mean waiting time

$$w = \frac{1 - c}{1 - p} \left[ \frac{E(q_x)}{p} + \frac{1}{2c} \right] \quad (9)$$

where  $c = g/t > p$ .

**Lemma 1 (Convexity of Delay):** The mean waiting time  $w$  in (9) is convex in the arrival rate  $p$  and dummy rate  $d = c - p$ , although not jointly convex, for  $p \in [0, 1]$ ,  $c \in (p, 1]$ .

*Proof:* When  $c > 2p$  (so  $E(q_x) = 0$ ), the second derivatives of  $w$  with respect to  $p$  and  $c$  are

$$w_{pp} = \frac{1 - c}{c(1 - p)^3}, \quad w_{cc} = \frac{1}{c^3(1 - p)}, \quad w_{pc} = -\frac{1}{2c^2(1 - p)^2}$$

It can be seen that  $w_{pp} > 0$  and  $w_{cc} > 0$  since  $p, c \in [0, 1]$ . When  $c < 2p$  (so  $E(q_x) > 0$ ), we have

$$w_{pp} = (1 - c) \left[ \frac{1}{(c - p)^3} - \frac{1}{p^3} + \frac{1}{c(1 - p)^3} \right] \quad (10)$$

$$w_{cc} = \frac{1 - p}{(c - p)^3} + \frac{1}{c^3(1 - p)} \quad (11)$$

$$w_{pc} = -\frac{1}{2} \left[ \frac{2 - p - c}{(c - p)^3} + \frac{1}{p^2} + \frac{1}{c^2(1 - p)^2} \right] \quad (12)$$

Since  $c > p$  then  $w_{cc} > 0$ . Since  $c < 2p$ ,  $w_{pp} > 0$ . However the Hessian need not be positive semidefinite (for example when  $p = 0.5$  and  $c = 0.9$ ). Since  $w_{pp} > 0$  and  $w_{cc} > 0$ , the delay  $w$  is convex in  $p$  and  $c$  individually, but since the Hessian is not positive semidefinite  $w$  is not jointly convex in these variables. Now  $d = c - p$  is a linear function of  $c$  and  $p$ . Convexity is preserved under linear transformations and so the stated result follows. ■

Note that since the on time  $g$  and cycle time  $\tau$  are expressed as numbers of slots, they are integer valued  $g, \tau \in \mathbb{N}$  and so the domain of ratio  $c$  is the rational numbers  $\mathbb{Q}$  rather than the real-valued numbers. We therefore consider the relaxed problem where  $c$  takes values in  $(p, 1] \subset \mathbb{R}$  in order to ensure convexity. This relaxation does not, however, entail any loss of generality. Rather than using fixed on time  $g \in \mathbb{N}$ , define a sequence  $g_k \in \mathcal{N}$ ,  $k = 1, 2, \dots$ . For a given value of  $c \in \mathbb{R}$ ,

by queue continuity an integer-valued sequence  $g_k$  exists such  $|\sum_k g_k - \hat{g}|$  is bounded and so the waiting time can be made arbitrarily close to that with specified  $\hat{g} = c\tau \in \mathbb{R}$ .

### C. Proportional Fair Allocation

We consider the following optimisation  $P$  which maximises log-rate (and so is proportional fair) subject to network and traffic shaping delay constraints:

$$\min_{\mathbf{p}, \mathbf{d}} U(\mathbf{p}) := \sum_{f \in \mathcal{F}} -\log(p_f) \quad (13)$$

$$\text{s.t. } \mathbf{w} \preceq \boldsymbol{\sigma} \quad (14)$$

$$\sum_{f \in \mathcal{F}} \frac{p_f + d_f}{\psi_f} \leq 1 \quad (15)$$

$$\mathbf{p}, \mathbf{d} \in [0, 1]^{|\mathcal{F}|} \quad (16)$$

with

$$w_f = \frac{1 - (p_f + d_f)}{2(1 - p_f)} \left[ \max\left\{ \frac{(p_f - d_f)(1 - p_f)}{p_f d_f}, 0 \right\} + \frac{1}{(p_f + d_f)} \right] \quad (17)$$

for all flows  $f \in \mathcal{F}$  and where  $\mathbf{w} = [w_f]$ ,  $\boldsymbol{\sigma} = [\sigma_f]$ ,  $\mathbf{p} = [p_f]$ ,  $\mathbf{d} = [d_f]$  for  $f \in \mathcal{F}$  are vectors in  $\mathbb{R}^{|\mathcal{F}|}$ . Constraints (14) impose the requirement that flow delay deadlines are met, while (15) ensures that the flow rates (including both information and dummy transmissions) can be scheduled by the network.

### D. Solving Non-Convex Optimisation $P$

Optimisation  $P$  is not convex because, by Lemma 1, the delay constraints are not jointly convex in  $\mathbf{p}$  and  $\mathbf{d}$ . Nevertheless, these constraints are convex in  $\mathbf{p}$  and  $\mathbf{d}$  individually. This suggests the use of an alternating approach to solve  $P$ . Namely, solve for  $\mathbf{p}$  holding  $\mathbf{d}$  constant, then solve for  $\mathbf{d}$  holding  $\mathbf{p}$  constant. Let  $\mathbf{p}_k^*$ ,  $\mathbf{d}_k^*$ ,  $k = 1, 2, \dots$  denote the sequence of alternating solutions found in this way. Each solution is feasible for problem  $P$ . Further, since each individual optimisation is convex, we can find a global minimum and so  $U(\mathbf{p}_{k+1}^*) \leq U(\mathbf{p}_k^*)$ . Hence, the sequence  $\mathbf{p}_k^*$ ,  $\mathbf{d}_k^*$  is guaranteed to converge to a feasible stationary point of problem  $P$ . While this stationary point is, in general, sub-optimal, in practice we have found that it is usually close to a global optimum.

To carry out each optimisation we use a subgradient method for simplicity and because of its suitability for distributed implementation. Of course other methods might also be used. The resulting procedure is summarised as follows:

---

#### Algorithm 1 Alternating Solution Method

---

```

iterate  $s$ :
  iterate  $t$ :
     $\mathbf{p}(t+1) = \mathbf{p}(t) - \alpha \partial_{\mathbf{p}} L_{\mathbf{p}}(\mathbf{p}(t), \mathbf{d}(s-1), \boldsymbol{\lambda}_p(t))$ 
     $\boldsymbol{\lambda}_p(t+1) = [\boldsymbol{\lambda}_p(t) + \alpha \partial_{\boldsymbol{\lambda}} L_{\mathbf{p}}(\mathbf{p}(t), \mathbf{d}(s-1), \boldsymbol{\lambda}_p(t))]^+$ 
  loop
  iterate  $t$ :
     $\mathbf{d}(t+1) = \mathbf{d}(t) - \alpha \partial_{\mathbf{d}} L_{\mathbf{d}}(\mathbf{p}(s), \mathbf{d}(t), \boldsymbol{\lambda}_d(t))$ 
     $\boldsymbol{\lambda}_d(t+1) = [\boldsymbol{\lambda}_d(t) + \alpha \partial_{\boldsymbol{\lambda}} L_{\mathbf{d}}(\mathbf{p}(s), \mathbf{d}(t), \boldsymbol{\lambda}_d(t))]^+$ 
  loop
loop

```

---




Fig. 6: Illustrating optimal throughput and dummy rate for delay deadlines  $\sigma_2 = 10$  and  $\sigma_1 = \{5, \dots, 15\}$ .

with Lagrangians,

$$\begin{aligned}
 L_{\mathbf{p}}(\mathbf{p}, \mathbf{d}^*, \boldsymbol{\lambda}) &= \sum_{f \in \mathcal{F}} -\log(p_f) + \sum_{f \in \mathcal{F}} \lambda_{1,f}(w_f - \sigma_f) + \\
 &\quad \lambda_2((\sum_{f \in \mathcal{F}} p_f + d_f^*) - 1) + \\
 &\quad \sum_{f \in \mathcal{F}} \lambda_{3,f}(p_f - 1) - \sum_{f \in \mathcal{F}} \lambda_{4,f} p_f \\
 L_{\mathbf{d}}(\mathbf{p}^*, \mathbf{d}, \boldsymbol{\lambda}) &= \sum_{f \in \mathcal{F}} \lambda_{1,f}(w_f - \sigma_f) + \\
 &\quad \lambda_2((\sum_{f \in \mathcal{F}} p_f^* + d_f) - 1) - \sum_{f \in \mathcal{F}} \lambda_{5,f} d_f
 \end{aligned}$$

### E. Examples

We present a number of examples to illustrate the proportional fair allocation in a fully private shared network.

*Example 1: Fully private network.* Consider a network with two users having different delay deadlines. We let delay deadline  $\sigma_2 = 10$  and vary  $\sigma_1$  between  $\{\sigma_2 - 5, \sigma_2 + 5\}$  to observe the impact of the delay deadline on users' network share in a fully private network. Figure 6 shows the proportional fair  $p^*$  and  $d^*$  vs  $\sigma_1$ . It can be seen that users' throughput and dummy rate are proportional to their delay deadlines *i.e.* users with lower deadline are allowed to transmit more real and dummy traffic resulting a larger network share.

*Example 2: Mix of private and non-private flows.* As already noted, users in a shared network need to sacrifice throughput and/or delay to achieve full privacy. However, staying within our traffic shaping framework, a user can choose to ignore privacy by sending no dummy packets and using the full cycle length for transmitting information packets. Of course this allows an adversary to see their packet arrivals. The optimization problem for this scenario is similar to 13 except that the delay and dummy rate constraints are now

$$\begin{aligned}
 w_f &\leq \sigma_f, & f \in \mathcal{F}_{\text{private}} \\
 d_f &> 0, & f \in \mathcal{F}_{\text{private}} \\
 d_f &= 0, & f \in \mathcal{F} - \mathcal{F}_{\text{private}}
 \end{aligned}$$

where  $\mathcal{F}_{\text{private}} \subset \mathcal{F}$  is the set of private users.

In this example we consider similar conditions to those in Example 1, but now User 2 is non-private. Figure 7 shows



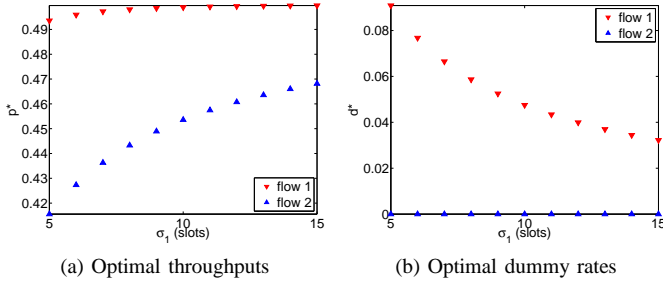


Fig. 7: Illustrating optimal throughput and dummy rate for a mix of private and non-private users. User 1 is private with  $\sigma_1 = \{5, \dots, 15\}$  while User 2 ignores privacy by transmitting no dummy packets and using the full cycle length for transmitting information packets.

$p^*$  and  $d^*$  vs the delay deadline of User 1. It can be seen, that User 1 consistently gets higher throughput than in a fully private network (compare Figures 7a and 6a).

## VI. SUMMARY AND CONCLUSIONS

In this paper, we introduce a rate allocation scheme for private shared networks. First, a defence is proposed against timing only traffic analysis attacks which protects the user by transforming their packet arrival time sequence into one which contains no information about the packet arrival pattern of the original sequence. The transformation however imposes a delay on transmission and consumes bandwidth by transmitting dummy traffic. We address a shared network scenario where the performance of one user can affect the network experience of another. This leads to a further analysis of the resulting trade-off between user privacy and quality of experience and to the design of a proportional fair rate allocation algorithm.

## REFERENCES

- [1] A. Banchs, P. Serrano, and H. Oliver. Proportional Fair Throughput Allocation in Multirate IEEE 802.11e Wireless LANs. *Wirel. Netw.*, 13(5):649–662, October 2007.
- [2] D. L. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM*, 24(2):84–90, February 1981.
- [3] J. N. Darroch et al. On The Traffic-Light Queue. *The Annals of Mathematical Statistics*, 35(1):380–388, 1964.
- [4] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton. Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 332–346, May 2012.
- [5] A. Eryilmaz and R. Srikant. Fair Resource Allocation in Wireless Networks Using Queue-length-based Scheduling and Congestion Control. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 3, pages 1794–1803. IEEE, 2005.
- [6] S. Feghhi and D. J. Leith. A Web Traffic Analysis Attack Using Only Timing Information. *arXiv*, abs/1410.2087, 2014.
- [7] S. Feghhi and D. J. Leith. A First-Hop Traffic Analysis Attack Against Femtocell. In *IEEE Consumer Communications and Networking Conference (CCNC)*, 2016.
- [8] Y. Guan, X. Fu, D. Xuan, P.U. Shenoy, R. Bettati, and W. Zhao. NetCam: Camouflaging Network Traffic for QoS-guaranteed Mission Critical Applications. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 31(4):253–265, Jul 2001.
- [9] S. Kadloor and N. Kiyavash. Delay Optimal Policies Offer Very Little Privacy. In *INFOCOM, 2013 Proceedings IEEE*, pages 2454–2462, April 2013.
- [10] J. F. C. Kingman. On Queues in Heavy Traffic. *Journal of the Royal Statistical Society. Series B (Methodological)*, 24(2):383–392, 1962.
- [11] L. Massoulié and J. Roberts. Bandwidth Sharing: Objectives and Algorithms. In *INFOCOM ’99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1395–1403 vol.3, Mar 1999.
- [12] A. Mishra and P. Venkitasubramaniam. Source Anonymity in Fair Scheduling: A Case for the Proportional Method. In *Communications (ICC), 2012 IEEE International Conference on*, pages 1118–1122. IEEE, 2012.
- [13] Abhishek Mishra. Tradeoffs Between Anonymity and Quality of Services in Data Networking and Signaling Games. 2015.
- [14] G. F. Newell. Queues for a Fixed-Cycle Traffic Light. *Ann. Math. Statist.*, 31(3):589–597, 09 1960.
- [15] M. S. van den Broek, J. S. H. van Leeuwen, I. J. B. F. Adan, and O. J. Boxma. Bounds and Approximations for the Fixed-Cycle Traffic-Light Queue. *Transportation Science*, 40(4):484–496, 2006.
- [16] A. M. White, A. R. Matthews, K. Z. Snow, and F. Monrose. Phonotactic Reconstruction of Encrypted VoIP Conversations: Hookt on Fon-iks. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 3–18, May 2011.
- [17] C. V. Wright, L. Ballard, S. E. Coull, F. Monrose, and G. M. Masson. Spot Me if You Can: Uncovering Spoken Phrases in Encrypted VoIP Conversations. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pages 35–49, May 2008.
- [18] T. Wu, Y. Xue, and Y. Cui. Preserving Traffic Privacy in Wireless Mesh Networks. In *Proceedings of the 2006 International Symposium on on World of Wireless, Mobile and Multimedia Networks, WOWMOM ’06*, pages 459–461, Washington, DC, USA, 2006. IEEE Computer Society.